

(21) Application No 0010711.0

(22) Date of Filing 03.05.2000

(71) Applicant(s)
3Com Corporation
(Incorporated in USA - Delaware)
5400 Bayfront Plaza, M/S 1308, Santa Clara,
California 95052-8145, United States of America

(72) Inventor(s)
Hamish D S Martin
Ronald Brown
Mark A Pearce

(74) Agent and/or Address for Service
Brookes Batchellor
102-108 Clerkenwell Road, LONDON, EC1M 5SA,
United Kingdom

(51) INT CL⁷
G06F 1/00 , H04L 12/22 29/06

(52) UK CL (Edition S)
H4P PEUX

(56) Documents Cited
WO 00/54458 A1 US 5991881 A US 5621889 A
<http://www.ticm.com/kb/faq/idsfaq.html> FAQ:
Network Intrusion Detection Systems
[http://www.iss.net/customer_care/resource_center/
product_lit/Real_SecureTM_Signatures, June 2000](http://www.iss.net/customer_care/resource_center/product_lit/Real_SecureTM_Signatures_June_2000)
[http://www.cisco.com/warp/public/cc/pd/sqsw/
sqidsz/tech/ntran_tc.htm](http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/tech/ntran_tc.htm)

(58) Field of Search
UK CL (Edition R) H4P PEUX PPA PPEB PPEC
INT CL⁷ G06F 1/00 17/27 17/30 17/60 , H04L 12/22
29/06
Online: EPODOC, JAPIO, WPI, Selected internet sites

(54) Abstract Title
Network intrusion detector which detects pre-attack probes or scans

(57) A network intrusion detection system (IDS) employs a method which looks at network traffic data and determines for each address external to the monitored network the number of communications or attempted communications with an address allocated to the monitored network. If this number exceeds a threshold within a predetermined period of time this may be indicative of an attempt to probe or scan the network and therefore the external address is flagged as a potential security threat or source of attack.

The system may also utilise a predetermined list of trusted external device addresses which the operator does not consider potential sources of attack and these addresses may be excluded from the method.

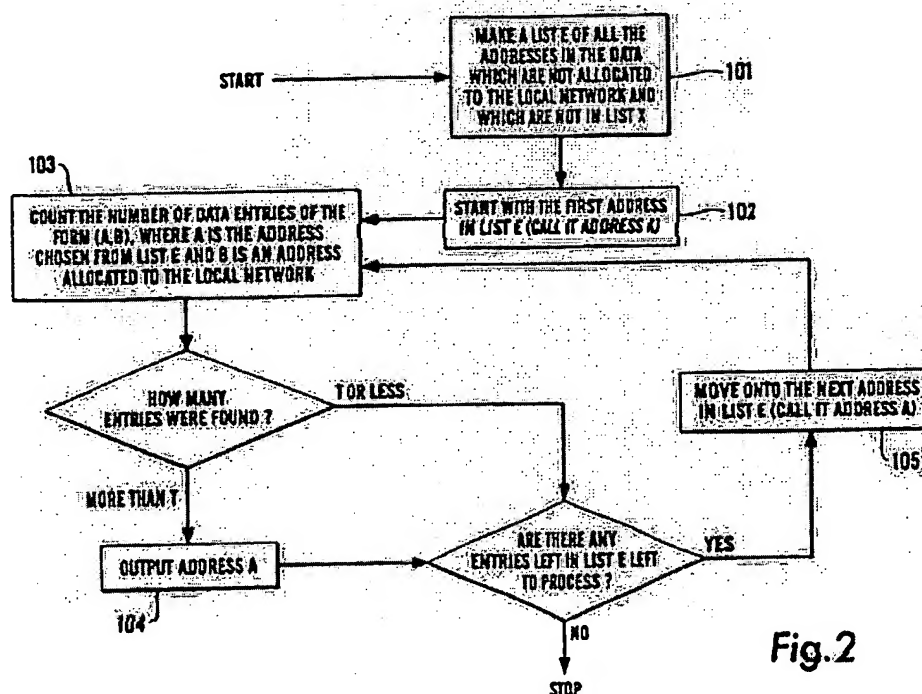


Fig.2




Fig. 1

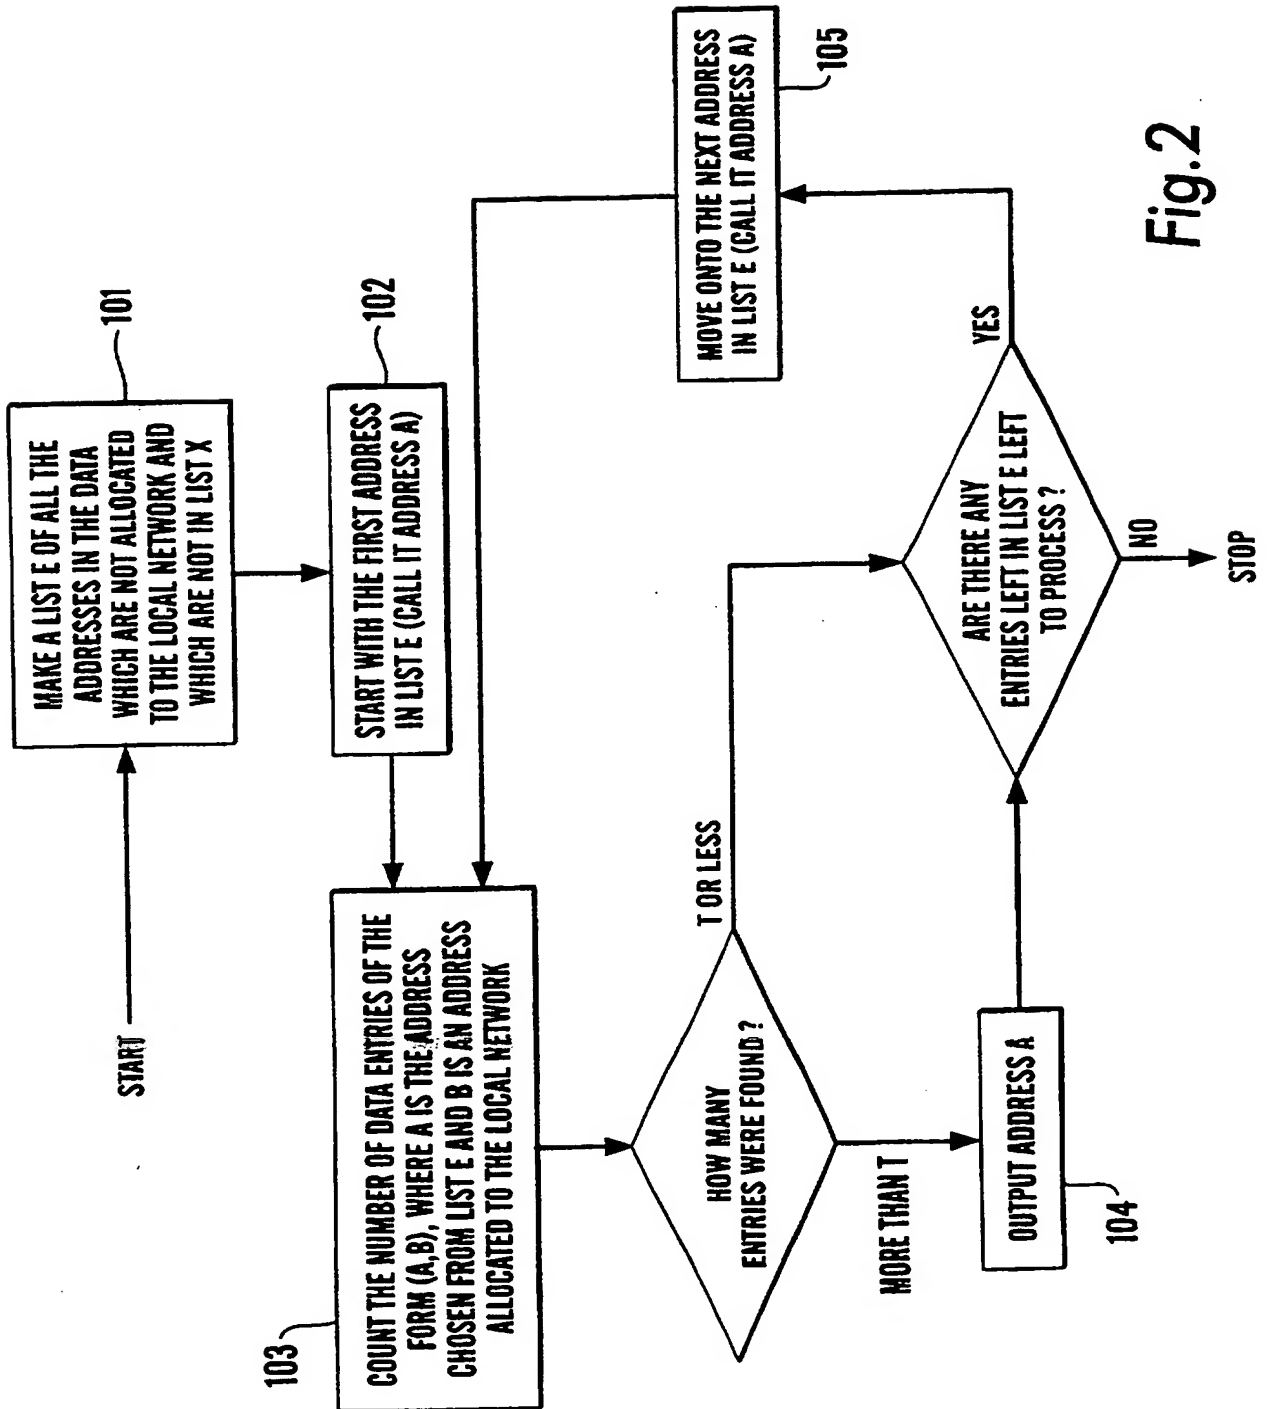


Fig. 2

DETECTION OF AN ATTACK SUCH AS A PRE-ATTACK ON A COMPUTER NETWORK

5 BACKGROUND OF THE INVENTION

The present invention relates to a method and apparatus for detection of an attack such as a pre-attack on a computer network by an unauthorised user.

10 Security is important to the manager of a modern computer network, be it a LAN (Local Area Network) or a WAN (Wide Area Network). Networks are usually attached to the Internet. Therefore, there is a constant risk that some malicious person from outside of a network may attempt to obtain access to the network and use this access to disrupt normal network activity or gain access to private information.

15 Many network managers use 'firewalls' (a device which filters traffic entering and leaving a computer network to protect it from malicious users) to protect their network from people outside the network. However, for many reasons firewalls are not suitable for all types of networks, since they may restrict the ability of legitimate
20 users to use the network and even where they are used, it is useful to have an additional level of security. We will describe a technique for detecting when someone from outside a network is attempting to access the network in an unauthorised way. The technique does not require a firewall in order to operate, and thus can be used as a complement to existing firewalls. For users who do not use firewalls, the method
25 described offers a way to try to detect unauthorised or malicious accesses to the network.

All devices on a network are identified by an 'address' (eg an IP address). When a
30 device wants to send data to another device, it typically marks the data with the destination address of the device it wants to communicate with and then puts this data onto the network, where is it forwarded to the correct device based on the destination address.

When a malicious person wishes to attack a network, it is usual for them to carry out what is referred to as a "pre-attack" on the network, that is to try to identify addresses which identify actual devices within the network. It would be useful to be able to deal with this problem.

5

Thus the arrangement of the invention allows the network to identify such a pre-attack.

SUMMARY OF THE INVENTION

10

The present invention provides a method for detecting a potential attack on a computer network comprising:

15 determining, over a period of time, and for each device address outside the local network, the number of communications or attempted communications with addresses within the local network, and where the number exceeds a predetermined number, identifying the external device address as a potential source of attack.

The present invention also provides a computer program on a computer readable medium or embodied in a carrier wave comprising the following steps:

20 from network traffic data which includes the source and destination addresses of traffic on the network, make a list E of all the addresses in the data which are not allocated to the local network and which are not in a list X;

choose a first address in list E;

25 count the number of data entries of the form (AB) where A is the address chosen from list E and B is any address allocated to the local network;

if the number of such data entries is more than T, output address A;

determine if there any entries in list E left to process;

if yes, move on to the next address in list E and repeat preceding three steps;

if no, stop.

30

BRIEF DESCRIPTION OF THE DRAWINGS

A preferred embodiment of the invention will now be described by way of example only and with reference to the accompanying drawings in which:

Figure 1 is a diagrammatic view of a network incorporating a preferred embodiment of the invention, and

Figure 2 is a flow chart of the steps of the preferred embodiment of the invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

A computer network will usually have a network manager who may set up, and control a computer network. The network manager will normally have his own network supervisor's workstation or computer.

Referring to Figure 1 there is shown a network comprising a plurality of devices in the form of a network supervisor's workstation or computer 11, other workstations 12B - E, hubs 13A, 13B, and switch 14. The network is a simple network and is set out for purposes of illustration only. Other configurations and arrangements, may be used.

The devices are connected together by means of links 16A - H which may be hard wired and utilise any desired protocol, and link 16F which is a wireless link.

The network supervisor's workstation includes, in addition to a visual display unit 18, a central processing unit or signal processor 19, a selector which may be in the form of a mouse 22, a program store 21 which may comprise, for example, a CD drive, a floppy disk drive or a zip drive, and a memory 17 for storing a program which may have been loaded from the program store 21 or downloaded for example via Internet from a website.

In a preferred arrangement, the computer 11 may, on command from the selector 22, process signals from the memory 17 by the signal processor 19 and provide on the visual display unit 18 a network map showing each of the devices and the links

therebetween. In the examples shown, the network is simple but of course in many instances the network will be considerably more complex and it may be necessary to arrange that the visual display unit 18 only shows a simplified version or only part of the network at any one time.

5

In order to initialise a computer network, a network manager (or the installer) needs to assign addresses to all the devices on the network. Typically, the manager of a computer network will receive an allocation of possible addresses for devices. The manager may assign any of these addresses to the devices that are actually attached to the network at his discretion. It is usual to have many more allocated addresses than devices on the network and typically only a small fraction of the allocated addresses are actually used by devices on the network.

10

In order to access devices on a network, the malicious user needs to find out the addresses of the devices on the network that can be illicitly accessed. Since network managers usually try to protect their devices against malicious users, a malicious user may only be able to gain access to a small number of the devices on the network, or possibly no devices at all.

15

In order to find out the addresses of devices that can be illicitly accessed, a user will typically perform a 'pre-attack probe'. Since the range of addresses allocated to a particular network are public knowledge, a 'pre-attack probe' involves using a program to test in turn every address within the range of addresses allocated to the network. For every address allocated to the network, the pre-attack program assumes that there is a device assigned to that address, and attempts to contact the device at the address to determine if the device is susceptible to illicit access. If there is a device allocated to that address, then the device will respond to the 'pre-attack probe' and the malicious user can then attempt to access this device for his own illicit purposes. If there is no device allocated to the address, then there will be no response to the 'pre-attack probe'.

20

25

30

Since every device on the network must be using one of the allocated addresses, if allowed to continue access, in time the 'pre-attack probe' is certain to discover all the devices that are susceptible to illicit access.

5 So called 'Pre-attack probes' are a common precursor to other types of illicit network access. We will now describe a way to detect malicious users by detecting 'pre-attack probes' using information collected from the network.

10 As is well known, the network manager will normally have installed on his workstation a program that enables him to understand the technical operation of the network. Some of the devices within the network will be "managed" devices, that is devices which include a so called "agent" which collects and stores data relating to the operation of that device and the traffic passing through, to, or from it. The network manager's computer, using the relevant software, interrogates the agent of
15 each device using a known protocol, such as SNMP (Simple Network Management Protocol).

20 A typical way of arranging this is to use a device called a RMON or RMON2 (Remote Monitoring Specification) probe for collecting data about the activity of devices on the network. Such a device uses the SNMP to transfer RMON or RMON2 information to the network management computer. The RMON2 standard is defined in IETF RFC2021. The network manager's computer includes a store (memory) in which the traffic information is stored for a period of time (day, week, month). This historical information may be stored as a database.

25 This traffic information may be provided to the network manager in any convenient form such as a table or graphic on his VDU (Visual Display Unit).

30 In accordance with the preferred of the arrangement, to detect 'pre-attack probes' on a particular network (which we will call the *local network*), data is required which records the pattern of traffic on the network over a selected time interval. This data must record every pair of addresses A and B where a device using address A has

attempted to communicate with a device using address B during the time period. A highly suitable source of data would be the RMON2 probe (in RMON2 either the allMatrix or nlMatrix tables could provide this information).

5 The data must include every communication between an address outside the local network and an address within the local network. Ideally, the data would also include attempts to communicate with addresses which are allocated to the local network but which have no device associated with them. This data is most sensibly collected by monitoring traffic flowing between the local network and the outside world.

10 The data is collected by a network management computer over the course of a few minutes or hours and then the following analysis is performed on the data by means of an algorithm (program), and the network manager alerted by the application if a 'pre-attack probe' is detected. This would often mean that the network manager would be
15 informed of the 'pre-attack probe' before the malicious user had a chance to do anything bad. This process is repeated continuously to provide constant monitoring for 'pre-attack probes.'

20 The preferred method of the invention is carried out under the control of the network manager's work station or computer and in particular by means of a program controlling the process of that computer or elsewhere in the system.

25 The program for controlling the operation of the invention may be provided on a computer readable medium, such as a CD, or a floppy disk, or a zip drive disk carrying the program or its equivalent, or may be provided on a computer or computer memory carrying the website of, for example, the supplier of the network products. The program may be downloaded from which ever appropriate source and used to control the processor to carry out the steps of the invention as described.

30 For the purpose of explanation it is assumed that the traffic data is presented as a list E of address pairs of the form (A,B). An entry in the list E of (A,B) indicates that a device with address A communicated or attempted to communicate with a device with address B during the time interval which the data corresponds to. (There may,

however, be no device with address B). Although the precise representation of addresses in the list is not important, devices will usually be represented by IP address, since this is the form in which RMON2 data is presented. If an RMON probe is used it normally provides the traffic data in the form of an 'nlMatrix' table of data.

5 The nlMatrix table records information about all conversations between devices, and stores them, using absolute counters, in a table which is ordered by network-layer protocol (e.g. 'IP') and source and destination addresses (i.e. IP addresses). An agent would record an entry in this table for every conversation which it has 'seen'.

10 Typically, a row in the stored table would look like this:

Protocol	Source Address	Destination Address	Packets	Bytes
IP	123.45.67.89	98.76.54.32	400	50000

The 'packets' and 'bytes' columns represent how many packets and how many bytes have been seen en route, travelling from the source address to the destination address.

15 The precise details of the table are given in the [RMON2 RFC].

The following is a simple example of how this RMON2 data could be used to form data into a list, which would be in a suitable form for the algorithm:

- 20
1. Create an empty list L.
 2. A program would read the entire nlMatrix table from a number of RMON2 agents.
 3. The program would wait for a predetermined period of time (e.g. 10 minutes).
 4. The program would read the entire nlMatrix tables from the same agents again.
 5. If the 'packets' count has changed on any row between step 1 and step 3, then an
- 25 entry is added to list L containing (source address, destination address).
6. Remove duplicate entries from list L.

Thus list L contains a series of entries (A,B) which represent traffic flowing from A to B, as required by the algorithm.

There are a variety of other, more sophisticated techniques which could be used for generating data in the required form from RMON2 data. For example, the algorithm could use data that has been collected and stored in a historical database.

5 The algorithm assumes that it is possible to determine if an address is allocated to the local network. There are multiple ways to achieve this, including allowing the network manager to indicate which addresses are allocated to the network and/or discovering addresses of local devices using some automatic process referred to briefly above.

10 In essence, the algorithm provides a method for determining a pre-attack on the computer network comprising determining, over a period of time, and for each device address outside the local network, the number of communications or attempted communications with devices within the local network, and where the number exceeds a predetermined number (T), providing an indication that the external device
15 address is a potential source of pre-attack. This is the process at its simplest.

However, other uses of the network may result in traffic patterns that resemble the traffic pattern of a 'pre-attack probe'. For this reason, the algorithm below provides a list of trusted devices, called list X. This list is intended to contain a list of devices
20 which the network manager does not believe will ever be the source of a 'pre-attack probe' (for example, addresses in another network of the same company). The algorithm will never indicate that any trusted device from this list has performed a 'pre-attack probe' and so the likelihood of a false alarm is reduced.

25 The algorithm assumes the existence of a numerical threshold value T. T is the maximum number of devices on the local network that a legitimate device outside of the local network might access in the time period of data collected. The actual value used for T is based upon the length of time over which the data was collected and the size of the local network. Too low a value of T will make the algorithm produce
30 'false alarms', while too high a value may mean that genuine 'pre-attack probes' are not detected. A value of T might be 10 – 100, typically 50.

The output of the algorithm is a list of addresses of devices that appear to have attempted a 'pre-attack probe' during the time interval of the data collected.

The program may include an algorithm of the form set out in the Figure.

5

Thus the program may include the following steps:-

program step 101, from network traffic data which includes the source and destination addresses of traffic on the network, make a list E of all the (source) addresses in the data which are not allocated to the local network and which are not in list X;

10

program step 102, start with first address in list E (call it address A);

program step 103, count the number of data entries of the form (AB) where A is the address chosen from list E and B is any address allocated to the local network;

if the number of such data entries is more than T, in program step 104, output address A;

15

are there any entries in list E left to process?;

if no, stop;

if yes, at program step 105 move on to the next address in list E (call it address A) and return to program step 103.

20

The address or addresses outputted at step 104 will be passed to the network manager's computer and highlighted as the address(es) of a potential malicious person attempting to gain access to the computer and using a pre-attack technique.

25

The preferred method of the invention is carried out under the control of the network manager's work station or computer and in particular by means of a program controlling the process of that computer or elsewhere in the system.

30

The program for controlling the operation of the invention may be provided on a computer readable medium, such as a CD, or a floppy disk, or a zip drive disk carrying the program or its equivalent, or may be provided on a computer or computer memory carrying the website of, for example, the supplier of the network products.

The program may be downloaded from which ever appropriate source and used to control the processor to carry out the steps of the invention as described.

5 Note that the operation of the invention is not affected by the presence or absence of a firewall.

10 It should be noted that the arrangement described uses historical data. In other words, traffic data is collected over a period of time and is then analysed subsequently. In this way patterns and relationships that build up over a course of time can be readily identified. If one were to try to carry out the same process in real time, and to maintain the data for a short period of time, it can be difficult to establish patterns of use. Thus, for example, the pre-attack program might access a succession of addresses in the local network over the course of a period of time which may range from minutes through to an hour or two and by using historical data the relevant
15 information can be readily accessed and analysed and becomes apparent.

We have described the method in terms of detecting an attack such as a pre-attack. However the same technique can be used to determine other forms of pattern of usage.
20

The invention is not restricted to the details of the foregoing example.

CLAIMS

1. A method for detecting a potential attack on a computer network comprising:
determining, over a period of time, and for each device address outside the
5 local network, the number of communications or attempted communications with
addresses within the local network, and where the number exceeds a predetermined
number, identifying the external device address as a potential source of attack.
2. A method as claimed in claim 1 in which external device addresses on a
10 predetermined list are not indicated as potential sources of attack.
3. A method as claimed in claim 1 in which only some of the addresses within
the local area network are connected to devices.
- 15 4. A method as claimed in claim 1 including the steps of from network traffic
data which includes the source and destination addresses of traffic on the network,
making a list of all of the addresses of said devices outside the local area network,
and, starting with the first address in the list, counting the number of data entries
which include A and B and which represent network traffic passing between a source
20 address A chosen from the list and a destination address B allocated to the local
network, and, if the number of such data entries is more than a predetermined number
T, indicate that the address is a potential source of attack, and returning to the list and
processing the next address in the list until all of the addresses in the list have been
processed.
- 25 5. A method as claimed in claim 4 in which the network traffic data is historical.
6. A method as claimed in claim 5 in which the network traffic data is collected
by an RMON probe.

30

7. A method as claimed in claim 1 including excluding from access to the local area network the or each external device address identified as a potential source of attack.

5 8. A computer program on a computer readable medium for carrying out the method of claim 1.

9. A computer program on a computer readable medium comprising the following steps:

10 from network traffic data which includes the source and destination addresses of traffic on the network, make a list E of all the source addresses in the data which are not allocated to the local network and which are not in a list X;

choose a first address in list E;

15 count the number of data entries which include A and B and which represent network traffic passing between a source address A chosen from list E and a destination address B allocated to the local network;

if the number of such data entries is more than T, output address A;

determine if there any entries in list E left to process;

if yes, move on to the next address in list E and repeat preceding three steps;

20 if no, stop.

10. A computer program embodied in a carrier wave for carrying out the method of claim 1.

25 11. A computer program embodied in a carrier wave comprising the following steps:

make a list E of all the addresses in the data which are not allocated to the local network and which are not in a list X;

choose a first address in list E;

30 count the number of data entries of the form (AB) where A is the address chosen from list E and B is any address allocated to the local network;

if the number of such data entries is more than T, output address A;

if the number of such data entries is less than T , determine if there any entries in list E left to process;

if yes, move on to the next address in list E and repeat preceding four steps;

if no, stop.

5



INVESTOR IN PEOPLE

Application No: GB 0010711.0
Claims searched: 1-11

14

Examiner: Owen Wheeler
Date of search: 23 October 2000

Patents Act 1977 Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK CI (Ed.R): H4P (PEUX, PPA, PPEB, PPEC)

Int CI (Ed.7): G06F: 1/00, 17/27, 17/30, 17/60; H04L: 12/22, 29/06

Other: Online: EPODOC, JAPIO, WPI, Selected Internet Sites

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	WO 00/54458 A1 [PSIONIC SOFTWARE]	
A	US 5991881 A [CONKLIN]	
A	US 5621889 A [LERMUZEAUX]	
A	Robert Graham, FAQ: Network Intrusion Detection Systems, Version 0.8.3, published 21 March 2000, available from http://www.ticm.com/kb/faq/idsfaq.html In particular sections 1.6, 1.7, 1.9, 2.1	
A	Real Secure™ Signatures Version 5.0, June 2000, pages i-x, 15-17, 186-187 and 263-264 only, available from http://www.iss.net/customer_care/resource_center/product_lit/ In particular "Portscan Detection" and "UDP Port Scan".	
A	Technical Brief - Cisco Secure Intrusion Detection System Technical Overview, 3 Jul 2000, available from http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/tech/ntran_tc.htm In particular "Sensor Capabilities" pages 2-4.	

X Document indicating lack of novelty or inventive step
Y Document indicating lack of inventive step if combined with one or more other documents of same category.
& Member of the same patent family

A Document indicating technological background and/or state of the art.
P Document published on or after the declared priority date but before the filing date of this invention.
E Patent document published on or after, but with priority date earlier than, the filing date of this application.

THIS PAGE BLANK (USPTO)